

# Effects of Target Arrival Rate on Mission Performance of Cooperatively Controlled UAVs with Communication Constraints\*

Jason W. Mitchell<sup>†</sup>, Steven J. Rasmussen<sup>‡</sup>, and Andrew G. Sparks<sup>§</sup>

**Abstract**— We investigate the effects of target arrival rate on the communication and mission performance of cooperatively controlled uninhabited aerial vehicles with task allocation performed by iterative network flow. Specifically, we quantify the effect of arrival rate on observed statistics of communication and mission performance. The statistics of interest are peak communication data rate, execution defects. The effects are seen in a series of vehicle-target scenarios simulated in the U.S. Air Force Research Laboratory's MultiUAV environment.

## I. INTRODUCTION

Cooperative control of uninhabited aerial vehicles (UAVs) continues to be of significant interest as a means to improve mission performance. The ability to communicate information necessary to cooperatively execute the tasks of search, detect, classify, attack, and verify is an essential element in enhancing team performance.

While communication, in an ideal sense, may improve team performance in general, it is unlikely to be true for all cases of a real system. This is due to the typical control algorithm design practice of initially ignoring communication issues to reduce complexity. Thus, it is necessary to study the impact a posteriori.

Previous work [1–3] investigating communication effects on cooperative control has focused on target distributions that were uniformly distributed spatially within an area of responsibility (AoR). In this work, we investigate communication effects arising from a temporal target distribution within the AoR.

## II. BACKGROUND

We begin with a short description of MultiUAV2 including the general mission scenario architecture. This is followed by a brief summary of the data common to each mission scenario.

### A. Simulation Framework

The MultiUAV2 simulation package [4], [5] is capable of simulating multiple uninhabited aerial vehicles (UAV) that cooperate to accomplish a predefined mission. Individually,

the vehicles are capable of searching for, recognizing, attacking, and verifying the destruction of targets.

The purpose of the simulator is to provide an environment in which researchers can implement and analyze cooperative control algorithms. MultiUAV2 is built using a hierarchical decomposition where inter-vehicle communication is explicitly modeled. The package includes visualization tools and provides links to external programs for post-processing analysis. Each of the vehicle simulations include six-degree-of-freedom dynamics and embedded flight software (EFS). The EFS consists of a collection of *managers* or *agents* that control situational awareness and responses of the vehicles. In addition, the vehicle model includes an autopilot that provides waypoint navigation capability. The individual managers contained within the vehicles include: Tactical Maneuvering, Sensor, Target, Cooperation, Route, and Weapons. At the top level, these managers are coded as SIMULINK models, with supporting code written in both MATLAB script and C++.

### B. General Scenario

Let us consider a set of  $N$  simultaneously deployed vehicles indexed by  $i \in \mathbb{Z}^+[1, N]$ . The targets, which may be found by searching, are categorized according to the value associated with their destruction. The individual targets are indexed by  $j$  as they are found, so that we find  $j \in \mathbb{Z}^+[1, M]$  with  $V_j$  as the value of target  $j$ . The vehicles are provided no precise a priori information about the total number of targets or their initial locations. This information can only be obtained by the vehicles searching for and finding potential targets via Automatic Target Recognition (ATR) methodologies. The ATR process is modeled using a system that provides a probability that the target has been correctly classified. The probability of a successful classification is based on the viewing angle of the vehicle relative to the target, Rasmussen et al. [5]. For this exercise, the possibility of incorrect identification is not modeled, however targets are not attacked unless a 90% probability of correct identification is estimated. Further details of the ATR methodology can be found in Chandler and Pachter [6], with a detailed discussion available in Chandler and Pachter [7]. Once successfully classified as a target, an attack vehicle is selected. Upon reaching the selected target, this vehicle releases its munition and is subsequently declared an unavailable asset, i.e. attack is a terminal task. Finally, the selected target must be verified as destroyed to complete the target specific task chain.

Throughout the simulation, at each target state change

\*This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

<sup>†</sup>Aerospace Scientist, Emergent Space Technologies, Inc., 6301 Ivy Lane, Suite 700, Greenbelt, MD 20770-6334, Jason.Mitchell@emergentspace.com.

<sup>‡</sup>Senior Aerospace Engineer, General Dynamics, Advanced Information Engineering Services, Wright-Patterson AFB, OH 45433-7531, Steve.Rasmussen@wpafb.af.mil.

<sup>§</sup>Senior Control Systems Scientist, Control Sciences Division, Air Force Research Laboratory, Wright-Patterson AFB, OH 45433-7531, Andrew.Sparks@wpafb.af.mil.

# Report Documentation Page

Form Approved  
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE <b>DEC 2004</b>		2. REPORT TYPE		3. DATES COVERED <b>00-00-2004 to 00-00-2004</b>	
4. TITLE AND SUBTITLE <b>Effects of Target Arrival Rate on Mission Performance of Cooperatively Controlled UAVs with Communication Constraints</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Air Force Research Laboratory, Air Vehicles Directorate, Wright Patterson AFB, OH, 45433</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>					
13. SUPPLEMENTARY NOTES <b>The original document contains color images.</b>					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES <b>6</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

or task failure, a resource allocation algorithm is executed to compute task assignments. The resulting assignment is sub-optimal. However, Rasmussen et al. [8] has shown that these assignments are typically near-optimal in an average sense.

### C. Scenario Data

Individual scenarios are composed of five (5) vehicles and four (4) targets appearing within an AoR whose area is approximately  $64.8 \text{ km}^2$  ( $25.0 \text{ mi}^2$ ). The nominal individual vehicle properties are: constant velocity of  $112.8 \text{ m/s}$  ( $370.0 \text{ ft/s}$ ) or  $M \approx 0.33$ , constant altitude of  $205.7 \text{ m}$  ( $675.0 \text{ ft}$ ), minimum turn radius of  $609.6 \text{ m}$  ( $2000.0 \text{ ft}$ ), and sufficient fuel for the duration of the required operation. The sensor footprint is a rectangle with a leading edge width of  $600.0 \text{ m}$  ( $1968.5 \text{ ft}$ ) and a depth of  $250 \text{ m}$  ( $802.2 \text{ ft}$ ), and is offset  $875.0 \text{ m}$  ( $0.5 \text{ mi}$ ) to the fore of the vehicle.

Since effects associated with search are not the focus of this study, vehicles begin in a line formation, and initially follow preprogrammed waypoints for a *zamboni race* coverage pattern. The targets are spatially distributed within the AoR according to the arrival model described in the following section, with pose-angles selected uniformly randomly. A total of eighteen (18) sets of fifty (50) individual simulations [8] are replicated for each of three arrival models, for a total of 2,700 individual missions. The maximum mission time is  $t_f = 220 \text{ s}$  for each arrival model.

## III. MODELS

### A. Task Allocation: Iterative Network Flow

Task allocation is performed by an iterative network flow of depth three (3) [9–11]. The acquisition of new target state information triggers the formulation and solution of a new optimization problem that reflects a decision based on the team’s current relevant operating perception.

At each solution depth, it is not possible to simultaneously assign multiple vehicles to a single target, or multiple targets to a single vehicle due to the integrality property of the network flow. However, using the single assignment process iteratively, *tours* of multiple assignments can be generated [11]. This is done by solving the initial assignment problem once, and only finalizing the assignment with the shortest estimated arrival time. The problem can then be updated assuming that the allocated task is executed, updating target and vehicle states, and running the assignment again. This iteration can be repeated until all of the vehicles are assigned terminal tasks, or until the target assignment pool is fully exhausted. The target assignments are complete when classification, attack, and verification tasks have been assigned for all known targets.

### B. Communication

The effects based communication simulation used here is that found in Mitchell et al. [3], where messages are passed to simulate vehicle communication at each major model

update, which occurs at  $10 \text{ Hz}$ . Such a coarsely grained update is necessary to maintain a reasonable run-time for individual scenarios to complete on a desktop/personal computer, in a larger Monte-Carlo sense. The minor model update, which controls the vehicle dynamics and other underlying subsystems, is scheduled at  $100 \text{ Hz}$ .

A *broadcast* communication model is implicitly assumed for the vehicle communication. While not specifically targeted to address a particular physical implementation, such a model encompasses the typical view that the communications are time-division multiplexed.

1) *Data Rate*: As a consequence of the major model update, we define the *data rate* necessary at a given simulation step as the total size of the messages collected, in bits, divided by the duration of the model update, yielding a rate in bits/s. This simplistic definition is a result of the elementary requirement that each vehicle must have access to all the currently generated messages by the next major update in order to plan. Currently, all message data is represented in MATLAB using double-precision floating-point numbers, and in the computation of data rate, the message overhead is not considered, only the message payload. In a physical communication implementation there would be considerably more overhead, e.g. redundancy, error correction, encryption, etc. Thus, retaining double-precision in the ideal communication model remains a reasonable indicator of data rates, particularly since we are interested only in an initial estimate and, more importantly, a relative comparison of communication necessary to perform under various conditions.

2) *Information Requirements*: The implementation of the task allocation algorithms outlined above requires communication of information between vehicles. The resulting overall optimization problem can be characterized as both centralized and redundant, i.e. each vehicle computes its own network flow.

Momentarily disregarding communication issues, the problem, in general, requires a synchronized database of target and vehicle state information. With this, each vehicle computes the benefits for the arcs in the network described previously, and solves an optimization problem to maximize the total benefit. From Mitchell et al. [2], the MultiUAV2 network flow implementation requires the following communicated information: ATR data; target and vehicle positions; target, vehicle, and task status; and vehicle trajectory waypoints.

### C. Arrival Process

Target detection interarrival times are represented as a stationary Poisson process of constant rate  $\lambda$ . To increase the likelihood that targets will appear within the AoR, they are required to appear only after a minimum time offset, taken as  $t_{\min} = 30 \text{ s}$ . Our choices for the constant arrival rates  $\lambda_i \in \{2, 10, 20\} \text{ s}$  are motivated by a desire to detect all targets within the team’s initial pass into the AoR, which has a one-way transit time of  $t_{\text{TR}} \approx 70 \text{ s}$ .

Since both the number of events  $N_E$  that will occur within the AoR and the maximum time  $t_f$  are known as part of the simulation truth information set, we model the inter-arrival times as a sorted  $N_E$ -tuple of uniformly distributed random variates  $x_j \in U(0, 1)$  with  $j \in \mathbb{Z}^+[1, N_E]$  [12, p. 369]. Thus, the  $k^{\text{th}}$  event time is prescribed by

$$t_k = t_{\min} + \lfloor \lambda_i \sum_{j=0}^k x_j \rfloor, \quad k \leq N_E, \quad (1)$$

where  $\lfloor x \rfloor$  is the largest integer  $y$  such that  $y \leq x$ . The arrival times  $t_k$  are restricted to occur as integer multiples of the major model update.

1) *Assumptions:* Before proceeding further, we address the assumptions necessary to select this arrival process model. There are several modeling assumptions that must be considered here. First is the requirement of event serialization.

By selecting stationary Poisson process model, we enforce a sequence of single events separated by the modeled interarrival times. Previous work [3] using targets distributed spatially in a uniformly random fashion indicates that simultaneous detections occur, and inflict a worst case strain on communication network capacity. Clearly such detection event sequencing is not possible with the chosen model. This would seem to indicate that a better modeling choice would have been a *batch* or *compound* Poisson process [13]. However, the communication of the simultaneous detection of  $k$  targets is a subset of a larger group of simultaneous communication events that can be triggered by multiple simultaneous classification, attack or verification of targets in addition to multiple simultaneous task failures. Since there appears to be a sufficiently large set of these batch communication events, the current arrival model should not exclude such events entirely. An additional assumption to consider is the homogeneity of the selected arrival process model.

The stationary Poisson processes assumes that the number of events occurring in disjoint intervals are independent, i.e. possesses *independent increments*. It is unlikely that, once a target has been detected, an adversary would distribute assets both independently and identically. Considering the nature of integrated air defense systems, detection of additional assets would seem more likely after the discovery of certain specific asset types, e.g. early warning radar. This could create an arrival rush that would invalidate the assumption of homogeneity. In this case, a better choice may have been a nonstationary Poisson process that more realistically considered the correlation of interarrival times for individual target types. Unfortunately, this requires specific data to construct  $\lambda(t)$  that is generally security sensitive and so not typically available.

Defending the assumption of homogeneity brings into focus the primary motivation for selecting the current arrival process models: in the absence of specific data to support a particular arrival distribution, the stationary Poisson is a




Fig. 1. Maximum data rate cumulative average ratio  $\bar{y}_j / \bar{y}_{50}$ .

commonly preferred process model.

#### IV. TERMINATION ANALYSIS

When using Monte-Carlo methods to observe a performance measure, ensuring a sufficient statistical weight by performing an adequate number of replications is essential. In previous work [1–3], [8], we chose to look at the weight provided by the cumulative average,  $\bar{y}_j = \sum_{k=1}^j x_k$ , of the measure of interest. As in Ref. [3], our fundamental measure is the maximum data rate. Considering the first data set in Fig. 1, we see a slightly more useful form of the maximum data rate cumulative average, viz.  $\bar{y}_j / \bar{y}_{50}$ , than that used in Ref. [3]. This provides a clearer indication of  $\bar{y}_j$ 's relative weight for each  $\lambda_i$ . We notice from Fig. 1 that we remain within approximately 2% of the final cumulative average  $\bar{y}_{50}$  for  $\lambda_2, \lambda_3$  after only slightly less than 30 and 20 replications, respectively. However we must continue to nearly 40 replications to remain within 1% of  $\bar{y}_{50}$  for each. The cumulative average for  $\lambda_1$  is less well behaved, and begins to stabilize after 45 replications. Similar results are found in the remaining data sets.

In addition to the above, we would like to consider the communication data rate of each replication as a random variable  $X_j$ , and determine the relative dependence of  $X_i, X_j$  for  $i \neq j$  and  $i, j \in \mathbb{Z}^+[1, 50]$ . Our interest here is to verify that the data rates observed in the replications are, at worst, only weakly correlated, and thus sufficiently independent to be useful. Such a determination would lend credence to the notion that the maximum data rate could be treated as an independent and identically distributed (IID) random variable, since it is one particular observation of an independent random process.

The scenario data rate correlation coefficients  $\rho_{ij}$  from the first data set can be seen in Figs. 2–4. From this, we see that the data rate is weakly positively correlated for  $\lambda_2, \lambda_3$ , Figs. 3 and 4, with relatively few exceptions.




Fig. 2. Scenario data rate correlation coefficient:  $\rho_{ij}, i \neq j$ :  $\lambda_1 = 2$  s.

Although it is difficult to determine from the figures, the negatively correlated coefficients reside in the uncolored contour space and are valued such that  $\rho_{ij} < -0.01$ . The correlation coefficients for  $\lambda_1$  are seen in Fig. 2. Here, we see larger positive correlation, than for the  $\lambda_2, \lambda_3$  cases. However, given the comparatively small interval size, this is not unexpected. Fortunately, the increased positive correlation remains relatively small. As with the other cases, the negative correlations were  $\rho_{ij} < -0.01$ . Similar results are found in the remaining data sets.

In support of the IID nature of the maximum data rate, we perform a Kruskal-Wallis test of the maximum data rate over all data sets for each  $\lambda_i$ . From Table I, we see that the computed  $P$ -values are sufficiently far away from zero to establish process commonality. However, we note that  $P_{\lambda_1}$  is considerably smaller than  $P_{\lambda_2}, P_{\lambda_3}$ , suggesting a trend where in faster arrival rates may frustrate our desire for an IID random variable. Nevertheless, with the current data, we conclude that the scenario data rate maintains sufficient replication independence.

TABLE I  
KRUSKAL-WALLIS  $P$ -VALUES FOR MAXIMUM DATA RATE.

Arrival Rate	$P$ -value
$\lambda_1$	0.209
$\lambda_2$	0.448
$\lambda_3$	0.417

## V. RESULTS

To begin, we examine the number of *plan execution defects* encountered for each arrival rate  $\lambda_i$ . In a single data set, there are, nominally, 600 tasks to perform per arrival rate, given that there are a minimum of three (3) target states to process for each of four (4) targets in each of the fifty (50) replications. This is the minimum number of transitions that must occur since poor look-angles can cause




Fig. 3. Scenario data rate correlation coefficient:  $\rho_{ij}, i \neq j$ :  $\lambda_2 = 10$  s.




Fig. 4. Scenario data rate correlation coefficient:  $\rho_{ij}, i \neq j$ :  $\lambda_3 = 20$  s.

classification failures, and inability to meet timing windows may cause verification failures. In the case of *stolen* tasks, we imply that the plan did not react quickly enough to the changing environment. In the case of classification, as a result of vehicle positions, a vehicle is able to increase the ATR value of a target as a task that it was not assigned.

Fig. 5, presents a visual summary of the plan execution defects encountered. Here, we are interested in comparing the total number of defects occurring for a given arrival rate  $\lambda_i$ . For this, the notched and whiskered box-plots in Fig. 5 provide a visual  $t$ -test that allows for a straight forward comparison. Several trends are readily apparent. First, we notice that for all execution defects, the  $\lambda_1$  mean is measurably different, i.e. there is no notch overlap. The execution defect means for arrival rates  $\lambda_2, \lambda_3$  generally agree and are reasonably similar for all but failed classification, where we see a measurable increase as the arrival rate increases. A second trend is the relatively large number of




Fig. 5. Number of task execution defects for  $\lambda_i$ .




Fig. 6. Maximum data rate frequency distribution for  $\lambda_i$ .

stolen tasks associated with arrival rate  $\lambda_1$ . We reason that stolen tasks occur as targets arrive faster because there is increased target crowding compared to the sensor footprint size, so that it is more likely that a stray sensor will unintentionally increase a target ATR value or complete an unassigned verify. For failed and stolen verification, the increased arrival spacing provides better matching of flight route to planned route, thus decreasing missed timing window task failures. It is not immediately clear why failed classifies increase with increased arrival rate.

Histograms of the maximum data rate distributions are seen in Figs. 6. The distributions seen in Figs. 6b and 6c are similar to that found in [3], where targets were distributed spatially in a uniformly random manner, in that there appears to be two primary modes of operation: a median mode and a slightly more frequent lower quartile mode. In the cases of Figs. 6b and 6c, the lower portion of the




Fig. 7. Communication history: largest maximum data rate,  $\lambda_1 = 2$  s.




Fig. 8. Communication history: largest maximum data rate,  $\lambda_2 = 10$  s.

maximum data rate distribution comprises slightly more than 50% of the replications. This low maximum data rate corresponds to the ideal operational mode, where all communication events are serialized with relatively little sustained bursting. We also note from Fig. 6a that an operational mode more frequently occurs that requires a slightly higher maximum data rate. Thus, in general, as the arrival rate becomes faster, the maximum data rate is more likely to shift higher.

Figs. 7–9 represent the communication data rate history, of the first data set, containing each arrival rates'  $\lambda_i$  maximum data rate. We notice that as the arrival rate slows, the duration of communication bursts decreases and becomes more widely spread over the mission time, as expected. We note that the peak rate increases from  $\lambda_1$  and is nearly identical between  $\lambda_2$  and  $\lambda_3$ . It is surprising, however, that the least maximum data rate seen in Fig. 7 corresponds




Fig. 9. Communication history: largest maximum data rate,  $\lambda_3 = 20$  s.

to the fastest arrival rate,  $\lambda_1$ . To understand this, we must recall our previous statement that a faster arrival rate merely makes it statistically more likely that the maximum data rate will be higher; there is no guarantee.

## VI. CONCLUSIONS

In this work, the effect of the target arrival rate on the communication and mission performance of cooperatively controlled uninhabited aerial vehicles with task allocation performed by iterative network flow was investigated. In general, the mission performance with respect to plan execution defects was similar for each of the two slowest arrival rates, with the exception of failed classifications—the reason for the failed classification behaviour is as yet unclear. The fastest arrival rate did have a significant and measurable effect on plan execution defects, generally driving them higher, where as before, failed classification remained the exception. Regarding communication, in general, faster arrival rates drove the maximum communication data rate higher. It also caused significantly more burst communication as compared to the slower arrival rates.

Regarding the actual magnitudes of the maximum data rates, these should not be taken as exact requirements or

measures, particularly because no specific communication protocol or hardware implementation has been defined. Rather, the magnitudes should be seen to represent traditional engineering estimates that say more in their relative significance than individual significance. With that said, these values do indicate the amount of raw data necessary to drive the cooperative control algorithms, allowing for comparisons between individual implementations of an algorithm.

## REFERENCES

- [1] J. W. Mitchell, C. J. Schumacher, P. R. Chandler, and S. J. Rasmussen, "Communication delays in the cooperative control of wide area search munitions via iterative network flow," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2003.
- [2] J. W. Mitchell and A. G. Sparks, "Communication issues in the cooperative control of unmanned aerial vehicles," in *Proceedings of the Forty-First Annual Allerton Conference on Communication, Control, & Computing*, 2003.
- [3] J. W. Mitchell, S. J. Rasmussen, and A. G. Sparks, "Communication requirements in the cooperative control of wide area search munitions using iterative network flow," in *Proceedings of the Fourth International Conference on Cooperative Control and Optimization*, 2003.
- [4] S. J. Rasmussen and P. R. Chandler, "MultiUAV: A multiple UAV simulation for investigation of cooperative control," in *Proceedings of the Winter Simulation Conference*, 2002.
- [5] S. J. Rasmussen, J. W. Mitchell, C. Schulz, C. J. Schumacher, and P. R. Chandler, "A multiple UAV simulation for researchers," in *Proceedings of the AIAA Modeling and Simulation Technologies Conference*, 2003.
- [6] P. R. Chandler and M. N. Pachter, "UAV cooperative classification," in *Workshop on Cooperative Control and Optimization*. Kluwer Academic Publishers, 2001.
- [7] —, "Hierarchical control of autonomous teams," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2001.
- [8] S. J. Rasmussen, P. R. Chandler, J. W. Mitchell, C. J. Schumacher, and A. G. Sparks, "Optimal vs. heuristic assignment of cooperative autonomous unmanned air vehicles," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2003.
- [9] K. E. Nygard, P. R. Chandler, and M. Pachter, "Dynamic network flow optimization models for air vehicle resource allocation," in *Proceedings of the American Control Conference*, 2001.
- [10] C. J. Schumacher, P. R. Chandler, and S. J. Rasmussen, "Task allocation for wide area search munitions via network flow optimization," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2001.
- [11] —, "Task allocation for wide area search munitions via iterative network flow optimization," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2002.
- [12] A. M. Law and W. D. Kelton, *Simulation modeling and analysis*, 3rd ed. New York: McGraw-Hill Book Co., 2000, McGraw-Hill Series in Industrial Engineering and Management Science.
- [13] S. M. Ross, *Introduction to probability models*, 4th ed. Boston, MA: Academic Press Inc., 1989.